

# Written Communications and Suggestions from D Bayani to the Masters Advisory Committee to Date

David Bayani\*  
dcbayani@andrew.cmu.edu

October 5, 2020

## Contents

<b>1</b>	<b>Intent and Overview</b>	<b>2</b>
<b>2</b>	<b>Material Primarily In Regard To Understanding and Changing the Social Dynamics of SCS</b>	<b>3</b>
2.1	Initial Suggestions to MAC Relating to DEI Work . . . . .	3
2.1.1	Permission From Others to Share Content in Section 2.1 . . . . .	8
2.2	The Pipeline Leading Students to Join CMU as Graduate Students . . . . .	9
2.3	Clarification on Intent of “Pipeline” Discussion . . . . .	15
2.4	Followups ██████████: People for Overseeing the Social Analysis of SCS . . . .	16
<b>3</b>	<b>Briefly: A Few Points of Considerations for Discussing ML’s Impact and Actions CMU Can Take</b>	<b>18</b>
3.1	The Importance of Incorporating Awareness of Uses and Responsibility into the SCS Ethos .	18
3.1.1	Use of Examples or Case-Studies . . . . .	19
3.2	Educating While Attempting Not to Act as Ethics Police (If Deemed Desirable to Not Do So)	21
<b>4</b>	<b>In Regard to Archiving MAC Materials</b>	<b>23</b>
4.1	Suggestion for Next Iteration of Improving Archiving . . . . .	23
4.2	In Regard to Privacy and Consent to Sharing Other People’s Contributions . . . . .	23
4.3	Other Implementations Details and Thoughts . . . . .	25
4.3.1	Some High-Level Factors to Consider . . . . .	25
4.3.2	The CMU MAC Mail Server . . . . .	26
4.3.3	Backing Up Slack Content . . . . .	26
4.3.4	Google Drives . . . . .	27
4.3.5	Size Limits in GitHub . . . . .	27
4.4	Addendum on October 4th, 2020: Clarification of My Actions and Stance, and Highlighting Several Tradeoffs . . . . .	27
<b>5</b>	<b>File Size and Hash</b>	<b>29</b>

---

\* داود بياني

# 1 Intent and Overview

Contained herein are a number of the written communications from me (D Bayani) to the the Masters Advisory Committee (MAC) and other entities that are immediately relevant. Perhaps these will be useful for future reference or further consideration. In particular, as committee membership changes and time progresses, it would be good to gather and/or “soft archive” material<sup>1</sup> for those who wish to come up to speed on prior MAC work/ focuses. The Meeting Minutes certainly attempt to enable such retention and reference for the weekly meeting discussions, but the communications in this document fall outside that scope; it might be that only those who want to take a “deep dive” into MAC focuses who will care to read this. Whatever the case, here it is should it ever prove useful; one hopes it will not prove to be an embarrassment.

A note on scope: While this does not contain all written communications from me (for example, minor conversations on Slack, comments on an open letter to the SCS Dean, etc.), bundled here are the major, non-trivially sized communications produced thus far in my involvement with the MAC. While I would like to gather content of others for reference, it does not feel like my right to do so - nor do I have the resources (such as sufficient time) to tackle that further. Some aspects of this limitation do show up in this document, such as anonymization of other contributing parties in section 2.1, and thoughts on archiving<sup>2</sup> in section 4. With these constraints in place, I limit myself to bundling primarily material I have reign over, at the risk of appearing extremely hubrisic.

A take-away from this document: I likely would benefit from further discussions with humanities professors. Indeed, that wing of expertise and human-knowledge may be under-utilized here; not only the content of those fields, but the methods used to go about their business. For instance, my impression is that those areas (law, philosophy, cultural studies, etc.) have developed effective heuristics / strategies for discussing controversial, difficult, or abstract/vague subjects. Heuristics rarely turn out to be perfect — I think practitioners in humanities would be the first to say-so — but ignorance of prior work and re-inventing the wheel are rarely good things. Personally, I stopped myself several times while writing content for this document, realizing I was attempting to lay-out an amateurish re-invention. It is infeasible to expect most members of the CS community to become humanities experts, but it is possible to have an awareness and general appreciation of what those branches of human endeavor offer. This may be an obvious point given that much of this document’s content either discusses social-structure issues inside SCS or materials at the intersection of computer science and humanities. Obvious as it may be, I feel it is appropriate to be explicit at least once in acknowledging the value brought by expertise outside of CS while considering issues inside the CS School.

This writeup may have many deficits. I hope the reader knows enough to identify and understand all of them. I hope one day to learn enough to do the same.

—D Bayani,  
September 29<sup>th</sup>, 2020

---

<sup>1</sup>points of discussion, general information, etc.

<sup>2</sup>At the time of writing this, archiving is a subject that the committee is considering but has not yet made clear commitments to.

## **2 Material Primarily In Regard To Understanding and Changing the Social Dynamics of SCS**

### **2.1 Initial Suggestions to MAC Relating to DEI Work**

A note on anonymization of content in this section:

In the below printed emails, I have attempted to provide basic anonymization [REDACTED]. Details of my considerations are reported in section 2.1.1, and my thought process behind selecting the method of anonymization can be found in section 4.2 — noting that I still required permission since the committee has yet to adopt a policy covering privacy [REDACTED]. While these “privacy considerations” distract a bit from the primary DEI-related content of this section, I include them as an example of how these matters might be handled going forward (and perhaps setting an aspect of precedent).

[REDACTED]

D Bayani <dcbayani@andrew.cmu.edu>

---

## Question on Next Meeting and Where to Go For Further Details

7 messages

---

D Bayani <dcbayani@andrew.cmu.edu>

Tue, Jul 14, 2020 at 6:04 AM

To: masters-advisory-committee@cs.cmu.edu

To Whom It Concerns-

I was referred to this list from the PhD anti-racism committee, since it seems a more appropriate fit for me (given that I am a masters student). I would like to know when your next meeting is, and where I should go - I assume that, like the PhD working group, you have zoom meetings on a regular basis. Also, if you have a website you could point me to with more information, I would be thankful if you could provide the address.

-sincerely  
D Bayani

---

person1FirstName person1LastName <person1Email@andrew.cmu.edu>

Tue, Jul 14, 2020 at 8:48 PM

To: D Bayani <dcbayani@andrew.cmu.edu>

Cc: masters-advisory-committee@cs.cmu.edu

[REDACTED]

---

D Bayani <dcbayani@andrew.cmu.edu>

Thu, Jul 16, 2020 at 3:39 PM

To: person1FirstName person1LastName <person1Email@andrew.cmu.edu>

Thank you [REDACTED] for your timely reply, and my apologies for my delay.

I believe I would like to attend this meeting. In terms of what I would like to discuss and its length, it likely depends on the flow of the meeting up to that point. If I do end up raising any items (i.e., material beyond responding to other's suggestions), I would be surprised if I, personally, speak for any more than 7 minutes total (including in dialog), but how long that would extend the meeting due to other peoples' comments I cannot guess. If this is too vague for your tastes, I would be happy to try and collect some more concrete points.

[REDACTED] An immediate suggestion would be to semi-routinely advertise them via the SCS -MS mailing list(s) - as far as I am aware, that is not currently being done.

[Quoted text hidden]

---

person1FirstName person1LastName <person1Email@andrew.cmu.edu>

Thu, Jul 16, 2020 at 6:37 PM

To: D Bayani <dcbayani@andrew.cmu.edu>

Cc: masters-advisory-committee@cs.cmu.edu

[REDACTED]

person2FirstName person2LastName <person2Email@andrew.cmu.edu>  
To: person1FirstName person1LastName <person1Email@andrew.cmu.edu>  
Cc: D Bayani <dcbayani@andrew.cmu.edu>

Sat, Jul 18, 2020 at 11:55 AM

D Bayani <dcbayani@andrew.cmu.edu>  
To: person2FirstName person2LastName <person2Email@andrew.cmu.edu>  
Cc: person1FirstName person1LastName <person1Email@andrew.cmu.edu>

Sat, Jul 18, 2020 at 8:26 PM

As discussed today, I am following-up with material relevant to the suggestions I raised today and some of the pointers that were requested.

Below are some comments exchanged between myself and [redacted] in a similar vein. While I do not believe this conversation was private, I have not gotten explicit permission from them to share verbatim - as such, both to respect any desire for privacy and potentially reduce the length, I have edited the below to be shorter and remove references to other parties. Anything I edit (redaction, summary, or shortening, only) I will enclose in square brackets (i.e., "[ ]").

=====beginning content from other messages =====

**Suggestions for [the Committee]:**

Second, I wanted to share some concrete suggestions. I had shared with via email with [redacted] prior to the last meeting; I had asked that this be shared, but I need to follow-up about the status - I'm not complaining or anything, I'm just trying to do due-diligence to get the information across, not knowing how well integrated things are behind the scenes. [redacted]

"  
I list some example suggestions below, though I am not sure how practical it would be to carry them out given resources (such as time) available. In my discussion below, I primarily consider analysis of CMU itself, but similar discussions are applicable if the scope of [redacted] on the scale of multiple universities. Given that universities tend to be more structured than broader society (or at least relationships of interest tend to be more clear), I am not sure how well these suggestions would transfer beyond multi-university setting:

It might be worthwhile to have a quantitative look at what the university's demographic profile is now, not only in regard to counts based on self-reported information (which I believe the university already collects) but some attempt to capture homophily (in the network sense) based upon protected attributes. This analysis could be applied to the structure of roles and position-filling in CMU, as well as examination of how review boards are filled, etc. Another example of potential analysis, if CMU plans to attempt any interventions to address concerns, is the conducting of disparate impact analysis after intervention; this has a fair footprint in algorithmic fairness literature, for instance (e.g., <https://arxiv.org/abs/1703.00056>). Of course, this has a broader discussion that social sciences as a whole can take part in; they particularly could play a role in determining what to look for and what results mean.

This all said, there would seem to be at least two points against such an endeavour.

1. Since [redacted] the resources [redacted] doing deeper analysis may be beyond
2. Even if various sophisticated measures could be gathered, one has to be cautious using them since
  1. Looking at sophisticated systems through the lens of a few summary statistics often misses something worthwhile. In particular, trying to optimize one metric can tend to make other unmeasured but valuable things degrade.
  2. Interpreting statistics gathered requires non-trivial insight into the system under analysis, for example [https://en.wikipedia.org/wiki/Simpson%27s\\_paradox#UC\\_Berkeley\\_gender\\_bias](https://en.wikipedia.org/wiki/Simpson%27s_paradox#UC_Berkeley_gender_bias)

In light of the concerns above, particularly (1), perhaps one of the most productive things [redacted] could do - if [redacted] interested in what I have outlined here - is to encourage others, such as professors at CMU, to do this analysis. That is, have the people who are doing these sorts of diagnostics of other systems in the world to turn their lens inward on CMU.

One hopes that some at CMU might be doing this already and that their information is up-to-date; if so, leveraging that information - after having discussions with those who did the analysis to make sure it is understood properly in context - seems like a prudent thing to do. If it has been done in the past but not up to date, encouraging a revisiting may be good.

The social issues of concern can be sophisticated, and while simple counts can highlight concerning trends, they also certainly hide more nuanced aspects of the story. In particular, doing deeper dives gives a better picture of what illness CMU might exactly have, and thus give a better idea of how to treat them. Otherwise, we might only be able to tell if CMU has obvious symptoms of illness - missing what the illness might be exactly, and missing other conditions that might be asymptomatic but troubling.

Feel free to share my thoughts above with [redacted]. You can just copy and paste this indented section above, along with the introductory text prior ("I list some example suggestions below,..." to this paragraph). Feel free to share my andrew email with [redacted] in case they want to follow-up on any of these comments; I ask that it not be shared beyond that, however. As I said, I would be happy to stay in the loop regarding these discussions, but not sure how helpful I can be; I'll chime-in if a circumstance seems productive for me to speak-on or assist with.

Further - and please share this as well since I did not mention this to [redacted] - an example of how the network homophily could be conducted can be seen in

[https://link.springer.com/chapter/10.1007/978-3-319-57454-7\\_18](https://link.springer.com/chapter/10.1007/978-3-319-57454-7_18) . I am not saying that this exact approach should be used, but I want to give an example of how such an analysis could be done (measured, quantified, described, discovered, etc.) after getting the appropriate network structure and information (who-advises-who, who was on a review board, who is the employer of who, etc.) . What granularity to conduct such an analysis and what to include would have to be decided, but the point here is that it is possible and reasonably well defined after that.

[...]

[summary: there were some other items from my message and response from the other party. Most primarily involved logistics and clarifying priorities. For the most part those are not relevant to the information I am trying to convey here. However, some side-information discussed brought up tangential issues, which I raise below.

The next section of text contains a later email from the chain. Notice my text, followed by their response, followed by my-response-to-their-response]

[...]

**Some Points that I Would Like the Committee to Keep In Mind:** These are two items that come to my mind in review of the above emails that I would appreciate [redacted] keeping in mind. That is, while perhaps these are not immediately actionable, I believe that these are facts and perspectives that could (and perhaps should) be weighted during discussion and decision making. As such, these are not necessarily items to add to the agenda (though maybe you want to - I don't know), but thoughts to help shape the corners of the perspective held while pursuing this work. I list these here in part for reference, since I don't know how long I will be interacting with CMU or this specific [redacted] - either due to [redacted].

**A caution about demographic data:** Though this is likely of little surprise, it is worth reminding that games are often played with the demographic categories counted, reminiscent of gerrymandering. That is, there are groups that may be part of the collective that [redacted] is concerned about and yet counted as "white" or "non-white" when opportune. An example that comes to mind are many students of middle-eastern, north-african, or central-asian descent (full disclosure: I am among this collection)- a good story overviewing the issues are here: <https://www.latimes.com/projects/la-me-census-middle-east-north-africa-race/> . This community certainly has been impacted by policies that are of questionable intent, at least on the national scale - the travel ban really hit hard for some (<https://pittnews.com/article/116755/featured/academics-fight-immigration-ban/> , <https://www.cmu.edu/oie/news-and-events/archives/2017-2018/december/travel-ban-update.html>). There are of course issues beyond simply policy at play here, such as lived-experience, but I spare delving into that here due to its weight and complexity. This is just one collection of communities that may not be "appropriately" counted; I would be surprised if there were not others, particularly when the demographic data is (a) granular or (b) confuses self-reported skin-color (an important consideration in a society where it matters so much) with "race" or proxies there-of.

While I believe [redacted] is most focused on groups counted very clearly in typical demographic tallies - and rightfully so - I say this to give caution over taking all these counts at face value. In particular, this is not simply an issue of ensuring "everyone is counted" (an issue I suspect [redacted] is already aware of, at least speaking of demographics generally), but ensuring that "people are counted 'properly'". While perhaps the percentage of the CMU-population which this might be applicable to is small, this discussion helps elaborate what limitations exist in reporting, as well as "threats to validity" of reports.

[summary: They agreed with my point; clearly they are aware of the general issues involved in making these "definitions"/"distinctions"/"labels" for categories of people, and particularly that some major organizations do this poorly. From their point of view, reducing negative impacts experience by the BIPOC communities is the most effective/ efficient method of addressing issues here-raised since helping the worst-affected tends to help the majority of those communities negatively impacted.]

Noted and agreed.

**A particular measure to please keep in mind - homophily:** I strongly suggest [redacted] consider homophily (in the network sense) as an issue and/or driving factor. While I have discussed this as a quantitative idea earlier, here I simply mean it as a concept to be well aware of; I'm not saying [redacted] has to know how to calculate it, but know what it is and reflect on whether it is of non-trivial potential concern ([https://en.wikipedia.org/wiki/Network\\_homophily](https://en.wikipedia.org/wiki/Network_homophily)). It may be part of the factors giving CMU a sickness - a symptom as well as a cause. I refer to my email earlier in this chain for details to consider. Again, this is an item that I raise in hopes to *add to the perspective* [redacted] maintains in understanding and analyzing the situation, regardless of whether I personally am around to discuss it.

[summary: They agreed that homophily is important , and clearly knew what it was prior to my mentioning. It is not clear, however, if they meant it is important on the conceptual level or as matter immediately relevant to work in discussion. Essentially they asked that I expand on how I see this playing a concrete role in the work conducted. Note that in my response, I do not quote them - I raise rhetorical questions, not references to their questions.]

To address your question (though a bit indirectly):

Why is this relevant to making CMU a more equitable place for BIPOC ? Obviously this impacts the ease of gaining access to opportunities once on campus. Further, it impacts the admissions process : imagine the extreme scenario where a member of an admissions board intentionally preferences people whom they suspect have as many protected attributes similar to themselves as possible (even if certain protected attributes are not shown, it can be easy to find proxies). This extreme case, and particularly the intentional selection, seems unlikely - or at least we hope it is - but it is still cause for concern if research groups, admissions committees, of other gate-keepers / opportunity-givers at CMU behave in a way notably consistent with homophily regardless of whether it is intentional.

This we are already in agreement on - so why am I rehashing this? I have said this to clarify how I see it playing a role. At the moment, it is not clear to me that the BIPOC population in SCS is large enough to pick up much of a meaningful signal of homophily if it is occurring. But suppose it is, and showing itself by clustering together more specific shared-attributes among the non-BIPOC partition. What would happen if this behaviour - if present - was completely removed? Would the proportion of admitted BIPOC-population at CMU be more representative of the population pooled for students? Would members of the BIPOC community find campus to be a more inviting place? I suspect yes, and moreover I doubt that the opposite would occur in most reasonable situations.

My point is this: we are concerned in part about systemic discrimination and the campus/university culture that results in disparate impacts adversely impacting BIPOC communities. Homophily based on protected attributes (with granularity consistent with, but most likely finer-grain than BIPOC versus non-BIPOC), if



strong, is indicative of a possible cause or is another indicative symptom - a description of the systematic trends we most likely want to break. Notice that even if [REDACTED] decides nothing can be done to fight it, or not enough resources exist to do so, the behaviour can still be measured and monitored, in many cases even using public data. This in-itself gives a gauge of the university culture and possible improvements to it over time.

Measuring aside, what can be done about it if there is a problem? A variety of things. On one end are strategies similar to unconscious bias testing, which includes raising awareness in individuals. On the non-individual level, there could be monitoring and department reach-out to particularly bad and consistently behaving entities. Another possibility is if a multi-person deliberative body shows strong homophilic behaviour only when certain people are members, then there might be more incentive to no longer invite those people. There are a variety of approaches to consider, perhaps not too dissimilar to other styles of actions [REDACTED] is already considering - just applied to a different place.

Do I think this is an issue at CMU? I have good reason to suspect it is. Personally, I have seen behaviour and speech that are consistent with this at varying levels of explicitness, particularly in regard to social structures and research groups at CMU; one hopes these were misunderstandings or at most bad-apples, but such a position is a good way to put one's head in the sand. More importantly than my opinion, however, is that there are ways to gauge this. I am tempted to say "objective" ways to gauge this, but as I said in the original email here, there are choices people must make when setting up the analysis that make it non-trivial to determine what the "exclusive right-ways" to do it are. I suspect the best way to do it is to repeat the analysis several times (easy to do after the data is collected) with varying levels of attribute granularity.

=====ending content from other messages =====

#### Further Miscellaneous Comments and Clarifications:

- In the above, where I say "in many cases even using public data", I was referring to gathering the network structure. In most cases for most relevant attributes, the attribute data will not be available explicitly and publicly. Use of surveys coupled with some network imputation techniques might give reasonable guesses, but that delves further into technical details than I'd like to invoke here. Also, there is a host of literature on guessing protected attributes from proxies - such as gender from name - but evoking that possibility raises issues both technical and potentially ethical.
- Regarding the ProNe paper listed above: As I said, it is an example to give an idea of what the mechanics of the analysis might look like, not necessarily an indication that that exact thing should be used. There are some hiccups in that paper in regards to a few guarantees, but the methodology overall is still very informative and worth taking note of.
- A note in regards to who should look at these things: As I said [REDACTED] CMU has some of the brightest minds in regard to measuring social structures, definitely proving problems exist, and suggesting actions that actually lead to desired outcomes. It seems reasonable to *have some of the people who focus on this at CMU focus on CMU*. Indeed, as [REDACTED] was saying in today's meeting, the role of [REDACTED] is in large part to advocate for specific changes and resource allocations - [REDACTED] do not necessarily have the resources and expertise needed to carry out the details of the actions prescribed (analogy: just because you know you need a track from place X to place Y that can hold freight trains, it doesn't mean you yourself can build or know how to build the tracks). As discussed [REDACTED], however, there are some factors that impede this happening. For instance, the potential lack of publishability (due to lack of interest), and fears of various forms of reprisals - small carrot, big sticks. One hopes, however, that if the analysis done spans a few universities, then it both helps minimize the sharpness of any negative findings for CMU and increases publishability - make the sticks softer and make the carrot bigger; the downside here is that it makes the endeavor a much larger thing to bite into - likely large enough to choke the effort in the womb. This all said, I do also believe that there are knowledgeable professors at CMU who would - schedule providing- guide such an effort because it aligns with their values, without need for changing the incentive landscape.

-yours  
D Bayani

[Quoted text hidden]

D Bayani <dcbayani@andrew.cmu.edu>  
To: masters-advisory-committee@cs.cmu.edu

Sat, Jul 18, 2020 at 11:37 PM

I realized after sending that I should have sent the below to the whole committee.

As I say below, I make a number of these suggestions to help fill-out the corners of a well-developed perspective. Bear in mind these were in the context of discussion with [REDACTED], a few are from early interactions where [REDACTED] were simply fielding suggestions/ input from [REDACTED]. I caution against losing sight of the main goals and top-level Asks thus far composed by [REDACTED]. I wrote the below in the spirit of developing more of the path, not attempting to replace or divert it. I of course encourage each of the items I lay-out below - but only after the reader also examines the current proposals and Asks of [REDACTED] to provide the context these belong in.

----- Forwarded message -----

From: D Bayani <dcbayani@andrew.cmu.edu>  
Date: Sat, Jul 18, 2020 at 8:26 PM  
Subject: Re: Question on Next Meeting and Where to Go For Further Details  
To: person2FirstName person2LastName <person2Email@andrew.cmu.edu>  
Cc: person1FirstName person1LastName <person1Email@andrew.cmu.edu>

As discussed today, I am following-up with material relevant to the suggestions I raised today and some of the pointers that were requested.

Below are some comments exchanged between myself and [REDACTED] in a similar vein. While I do not believe this conversation was private, I have not gotten explicit permission from them to share verbatim - as such, both to respect any desire for privacy and potentially reduce the length, I have edited the below to be shorter and remove references to other parties. Anything I edit (redaction, summary, or shortening, only) I will enclose in square brackets (i.e., "[ ]").

=====beginning content from other messages =====

#### Suggestions for [the Committee]:

Second, I wanted to share some concrete suggestions.[...]

[Quoted text hidden]

[Quoted text hidden]

[Quoted text hidden]



D Bayani &lt;dcbayani@andrew.cmu.edu&gt;

## Question To Confirm It Is OK to Share Past Content with MAC

4 messages

D Bayani &lt;dcbayani@andrew.cmu.edu&gt;

Sat, Sep 19, 2020 at 3:58 PM

To: person1FirstName person1LastName &lt;person1Email@andrew.cmu.edu&gt;, person2Email\_2@gmail.com, person2Email\_1@andrew.cmu.edu

Hi person1FirstName, Hi person2FirstName,

I am gathering and bundling some of the past discussions for sharing with the MAC via the google drive. Below I have attached part of the content I hope to bundle: it is the initial long email chain for which I first interacted with the MAC and we discussed several ideas in non-trivial detail.

I would like to confirm that the content, as shown, is acceptable to both of you since you both commented in the original email chain. Notice your names and emails are substituted for with person<number>LastName, for instance. While the anonymization used is imperfect, I expect that members of the MAC act in good faith; further, it will be harder to infer as time progresses. Further, the MAC was CCed on most of these emails anyway. Personally, since both of you express balanced, factual information (as opposed to opinions or things suggestive of opinions), I did not see anything here to be concerned about - but, out of general respect for privacy, etc., I wanted to be sure you both are alright that this will be shared (bundled, without further content modification, with other, additional material).

I am of course willing to discuss if either of you have concerns or requests for further redactions. Otherwise, if you are both alright with this in the purview outlined here, a simple "yea, go ahead" suffices for me to proceed.

-yours

D Bayani

email14072020.pdf  
114K

person1FirstName person1LastName &lt;person1Email@andrew.cmu.edu&gt;

Sun, Sep 20, 2020 at 5:39 AM

To: D Bayani &lt;dcbayani@andrew.cmu.edu&gt;

Cc: person2Email\_2@gmail.com, person2FirstName person2LastName &lt;person2Email\_1@andrew.cmu.edu&gt;

person2FirstName person2LastName &lt;person2Email\_2@gmail.com&gt;

Sun, Sep 20, 2020 at 10:07 AM

To: person1FirstName person1LastName &lt;person1Email@andrew.cmu.edu&gt;

Cc: D Bayani &lt;dcbayani@andrew.cmu.edu&gt;

D Bayani &lt;dcbayani@andrew.cmu.edu&gt;

Sun, Sep 20, 2020 at 10:48 AM

To: person2FirstName person2LastName &lt;person2Email\_2@gmail.com&gt;

Cc: person1FirstName person1LastName &lt;person1Email@andrew.cmu.edu&gt;

Thank you both.

[Quoted text hidden]



## 2.2 The Pipeline Leading Students to Join CMU as Graduate Students



D Bayani <dcbayani@andrew.cmu.edu>

---

### DEI-Related Material : Regarding the Pipeline Leading to SCS Student Population and Available Actions Impacting It

---

D Bayani <dcbayani@andrew.cmu.edu>  
To: masters-advisory-committee@cs.cmu.edu

Sun, Jul 26, 2020 at 9:50 PM

Dear Committee Members,


Attached is a document I provided to the PhD Anti-Racism Working Group, which I am also providing here (1) in hopes it will be useful, (2) since I am transitioning more to working with the MS committee since I, myself, am an MS student not a PhD. I am not necessarily a fan of spamming the group, but I was previously told that providing suggestions via this email is acceptable ; I beg your pardon if I misunderstood.

Of the material here, I think the diagram at the end will prove the most useful, providing a common, concise reference to the parts of the process and maybe clarifying what information is being requested when this group contacts people (or maybe I am fooling myself... pardon my hubris either way). Please feel free to share the diagram with whoever you need to while pursuing this DEI work. I hope the writing proves more helpful than it is cumbersome.

-yours

D Bayani

---

 **scsStudentPipeline.pdf**  
264K

# Follow-up from CMU SCS Anti-Racism Working Group Meeting on 7/23/2020 : Regarding the Pipeline Leading to SCS Population

David Bayani  
School of Computer Science, Carnegie Mellon University  
dcbayani@andrew.cmu.edu

July 26, 2020

## 1 The Pipeline and Its Parts

Attached is a diagram of roughly what I believe the pipeline looks like that leads to a person becoming a member of CMU SCS. The diagram is primarily made of three different types of entities :

- Populations (blue): These are collections of people who have been selected-for up to that point in the pipeline.
- Processes (light purple): These are forces that, for the most-part, remove members of the population from (i) to (i+1). A subtlety is that from (4) to (5), new members might be added due to people transferring into SCS; this presumably accounts for only a small part of (5), at least on the graduate level (on the undergraduate level, however, it likely is a different story). This fine-point aside, the processes can mostly be thought of as filters.
- Monitors (red): This collection of entities sets standards for how the processes should conduct their business, monitors the populations at different steps of the pipeline, and has the ability to enact policies aimed at altering populations. Note that, for simplicity, the diagram shows monitors that observe each population and may be able to impact the process between them - however, in reality, some of these monitors may be missing, may only be able to effect some processes, or only observe parts of the pipeline. For example, there might only exist a monitor that has information on (0) and (4), and can only "actuate" some procedures in (P2). What actually exists in terms of monitors, their power, and their resources is something this working group is still attempting to discover, as far as I am aware.

A few points worth noting:

- As displayed, this pipeline is sufficiently general to consider undergraduate, masters and PhD students (separate pipelines for each of the three, but each pipeline would follow this schema).
- Obviously, I have adopted some terminology above to refer to relevant pieces of this pipeline, but please feel free to change this terminology - for example, using standard terms from control-theory. There is no need to re-invent the wheel; I should have named them in a way consistent with an established field of study to begin with.
- Expanding on (G2) in the diagram: As stated, (G1) by itself would simply check that (4) and (5) look similar in aggregate compared to the target, (0). Without examining the change of individual membership, however, this would allow (4) to lose many people during (P4) so long as by the end of (P4) a demographically similar set of people come in to fill the space generated. Whether or not this occurs or is likely to occur would have to be determined. This would clearly be an issue relevant to DEI if a handful of protected groups exit-and-enter SCS at a far higher rate than others. Such behaviour may be indicative of troubling aspects or areas needing improvement in (P4).

## 2 Examples of Interventions for Each Process in the Pipeline

During the meeting on 7/23/2020, where I discussed understanding this pipeline as the basis of questions I drafted, it was requested I follow-up with some suggestions on how parts of this pipeline might be actuated. That is, below is a sample of potential actions that can be performed by (M i) to effect (P i) if either the change in population between (i) and (i+1) or the difference between (i+1) and (0) run counter to our goals, (G1) and (G2) <sup>1</sup>. These are truly off the cuff - the primary intent in listing these is to begin brainstorming and demonstrate some available "levers to pull" for each process. I leave it to those more well versed in the dynamics of each process to suggest better actions and judge the likely effectiveness of proposed actions<sup>2</sup>.

- (P0): Effecting (P0) likely involves reaching out to people to encourage applying, as well as to boost people's exposure to "key" CS-material so to stoke interest. This is quite generic advise, I am aware. At the graduate level, this would seem to mean touching base with people currently in industry and senior undergraduates

<sup>1</sup>It is possible that changes early in the pipeline might only have noticeable impact on behaviour later on, in which case changing behaviour at earlier stages might be necessary even if measures appear fine there. Considering such down-stream phenomena, however, is the concern of an area expert. Short of having well-grounded expertise, proposing up-stream selection as culprits for behaviour changes between (i) and (i+1) would likely be idle speculation that could be used to favour any action equally well. It may be real, but you don't design a car starting with quantum physics, especially if you never studied the subject.

<sup>2</sup>Maybe this statement is a cop-out, but I would prefer to consult others than waste high-inertia resources on my under-informed interventions.

across multiple universities and multiple disciplines<sup>3</sup>. With regard to getting the message to individuals working in industry, this could be done by having professors - if they are willing - incorporate information into their talks or post-talks; note that this likely would attract people already involved in technology. Providing easier (though necessarily limited) access to some classes and / or creation of seminars/ tutorials to draw-in people might also help; these offerings could be curtailed to either those already working in technology or to those with less technical background. While the point of this document is primarily to discuss graduate-level education, note-worthy enrichment / outreach programs hosted by CMU are Andrew's Leap and C-Mites, which target high school and pre-high school students across South West Pennsylvania; while I do believe that these are excellent programs and likely increase applications to CS, it is not clear to me if they help (G1) and (G2).

- (P1): Similar to (P0) in that outreach would be involved - get people who want further education to apply to CMU SCS. Fee-wavers and waving the need for GREs might contribute.
- (P2): Alter how reviewers are selected, or - in the case that multiple people look at each application - ensure that the set of reviewers has variety in manners deemed relevant. Ensure the protected attributes and strong proxies for them are hidden when possible; this will likely be hard or have minimal effectiveness on its own since proxies for protected attributes may appear in critical sections of the applicant material, and wide sets of proxies exists. Further, in respect to attempting to hide protected attributes, additional resources would have to be allocated to the already resource-intense review process to ensure application-cleaning is done by a third-party prior to proceeding; some automation may be possible, but given the sophistication needed and the high-stakes of being wrong, automation would likely only provide trivial support.
- (P3): Again, outreach to those accepted. Given that people at this phase were members of (2) and (3), we might be allowed to know more about them individually and their values. Emphasizing that CMU is a welcoming and enjoyable place in a non-gimmicky way would be good... This can be done by personal phone-calls from relevant student representatives ; CMU already has phone-banks of students to call alumni to ask for money (CMU Telefund). I seem to recall such phone-calls exist for targeting new undergraduates, but, being a CMU-wide effort, the representative might not be from SCS; whether this still exists and the details of this program would have to be checked. In-person visits to campus also provide an opportunity for the CMU SCS community to show candidates that we *sincerely* share values with them.
- (P4): Better support for students from across a spectrum of experiences. Make sure people don't feel like outsiders. Better emphasis among SCS students that not everyone has the same background they do, and that aspects of their lives that

---

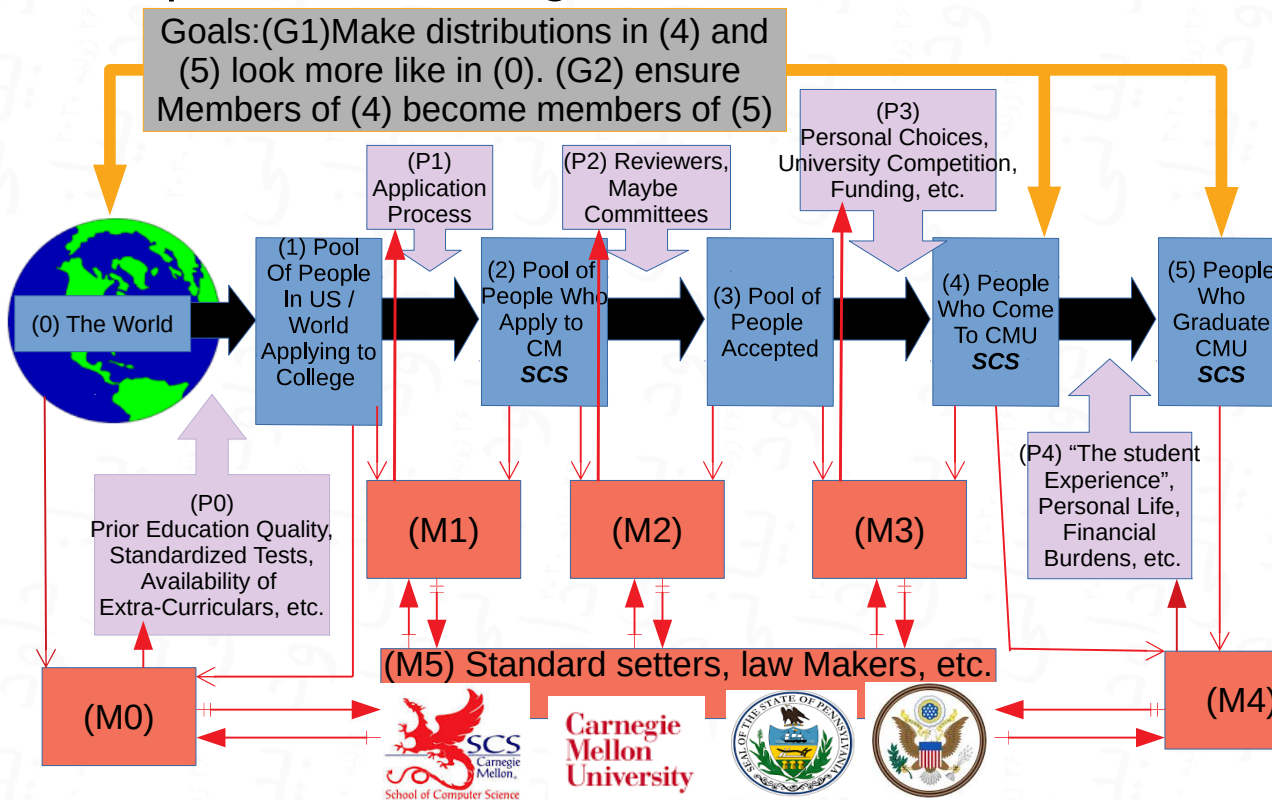
<sup>3</sup>Regarding multiple disciplines: A person who studied art or economics may like to pursue a further education in CS, for instance. CMU claims to pride itself on these cross-discipline stories.

they may have taken for granted are privileges that not all their classmates may have enjoyed: this extends to issues involving race, gender, income, prior educational quality, and home-life for instance. For some, resources to provide true community might not be available at CMU. Perhaps better collaboration with Pitt or other surrounding universities may help form a coalition of support that can assist students to muster-through CMU - but this suggestion sounds repetitive with existing student organizations, many of which already span multiple universities. One hopes that some potential changes exist which result in more people thriving as opposed to simply "mustering-through".

In the points above, several list outreach as an option. Providing such generic advice might make the discussion up until now seem under-whelming. The value of this perspective, however, is to know which part of the pipeline to focus outreach efforts on - not only does this impact resource provisioning, but it also impacts the type of outreach performed.

# Pipeline Producing Final SCS Student Distribution

Goals: (G1) Make distributions in (4) and (5) look more like in (0). (G2) ensure Members of (4) become members of (5)



## Key

A Population:  
Want summary  
Stats of  
distribution

Process:  
Shapes input  
Distribution to  
Output Distr.

Monitor for  
compliance

Summary  
stats.  
Standards  
Reports  
Actions

D Bayani July 25<sup>th</sup>, 2020



## 2.3 Clarification on Intent of “Pipeline” Discussion

To avoid misinterpretation, I want to be clear that my discussion of “the pipeline” above is **not** to suggest that CMU simply has “a supply-problem” and that “the problem would be solved if a wider, more varied pool of people applied”. This, in particular, seems to be the notion of a “pipeline” referred to in comments found in the public asks-letter to the SCS Dean. In contrast, the intent of this section’s content was to attempt to be more systematic when thinking about how we could attack the issue; at the time, I did not get the impression that discussion was sufficiently clear in regards to composition of the problem, what pieces would be impacted by proposed actions, or the relative effectiveness of potential actions.<sup>3</sup> It is possible that I missed avenues by which people become students or other factors (and perhaps it was hubristic for me to think I might help supply a schematic). At the very least, however, I hope some beneficial inspiration is drawn from this preliminary analysis of the problem and powers we have.

It is worth pointing-out that, in my off-the-cuff suggestions for actions, I did not simply suggest “attracting a wider range of ‘qualified’ students” as a panacea to SCS’s problems. Suggesting “supply-improvements” as a panacea appears to be the complaint behind references to “pipeline issues” in the open letter. A major practical concern with such a proposal is that CMU SCS could then cop-out and blame factors outside the university, while attempting only token actions inside the university — an unfortunate pattern that does not appear uncommon in large organizations. Practicality aside, the notion of a “supply-improvement panacea” is inherently troubling on its own in regard to truthfulness and honest institutional self-reflection.

The point here was to encourage more systematic thinking of what steps exist (as well as, implicitly, what steps might be added or removed) and contemplate actions we can take to change outcomes at each step. This was far from an attempt to shift blame outside of CMU SCS.

At the time of writing this, I have yet to receive any explicit complaints. Given the potential for confusion, however, I wanted to further obviate the distinctions.

---

<sup>3</sup> Some members of those discussions might already have had a break-down in mind and simply took it as granted in discussion. Even if so, however, I did not get the impression that everyone present had an organized or common view. Naturally, though, my impression might not have been fully accurate.

D Bayani <dcbayani@andrew.cmu.edu>

### Follow-Ups [REDACTED]

D Bayani <dcbayani@andrew.cmu.edu>

Sun, Aug 9, 2020 at 7:17 PM

It was requested [REDACTED] that I follow-up with some names of potential area-experts to consult for analysis of social patterns in SCS. This discussion was motivated in part from the conversation regarding potential analysis of homophily in the SCS community. I provide some notes/ suggestions on these fronts below ; I am happy to discuss any of these further. I of course give permission to have this forwarded to [REDACTED] if you think it is appropriate, but that is subject to your judgments.

#### Quick Aside: Illustration of How Homophily Can Be Quantified:

While the general concept I believe was clear, I do want to give an illustration of how, concretely, an analysis of homophily in SCS could be done (measured, quantified, described, discovered, etc.) after getting the appropriate network structure and information (who-advises-who, who was on a review board, who is the employer of who, etc.). One such illustration can be seen in [ProNe](#). What granularity to conduct such an analysis and what to include is another discussion, but the point here is that it is possible and reasonably well defined after that.

I am happy to discuss this point further if either of you are interested. I have provided written discussion and additional technical details to [REDACTED] on this issue, and would be happy to share that if you would like.

#### Suggestions of Relevant Personnel at CMU ( People Who Might Be Equipped and Willing to Guide SCS Social-Data Collection and Analysis Work):

The primary item from [REDACTED] meeting that motivated this email was a request for suggestions on personnel at CMU (not necessarily just SCS) who might be willing to oversee / conduct a social network analysis of the SCS community. I provide some such thoughts here.

#### Concretely: Some Names ( Personnel to Potentially Tap for this Endeavour):

To get to the material you probably most immediately wanted: some names of personnel potentially worth contacting. These are people that may be usefully for formally examining SCS's social structure and/or could suggest other members of the CMU community to also contact. I want to emphasize that latter part: even if these people are unwilling to do it themselves or you prefer someone else, given their knowledge of the CMU ecosystem for their areas, they could provide excellent suggestions for people (faculty, staff, and - as people who might actually do a lot of the leg-work - PhD students).

- Alexandra Chouldechova : She has a track record of using SCS to analyze issues relevant to DEI, and has deep knowledge of the sensitive issues in both the surface subject and the configurations of analysis. For instance, her work in algorithmic fairness (for example <https://arxiv.org/abs/1703.00056> ) would be good to draw on. I am not sure if she has a background in network analysis, but her expertise definitely would be valuable in setting up and drawing conclusions from any such analysis - for example, in choosing what attributes to examine, how they should be represented (categorical, binary, one-versus-all, etc.), their impacts on interpreting the results, potential biases that could creep in, and area-specific data processing / preparation issues.
- Ryan Ghani : he is an obvious asset in AI-and-social-good endeavours. Of particular interest here is his experience tackling real-world problems while working with policy makers and administrators; he understands what different stake-holders can do and how to effectively work with them. He is a mountain of expertise; here, I limit myself to only mentioning his knowledge of real-world data gathering-then-cleaning and the policy considerations he has been in-the-loop for.
- Kathleen M Carley - strong focus on social graph analysis of the flavor in discussion here. This includes both characterizing the behaviour of actors in a network and counter-actions that can be adopted to change the dynamics of a network. Indeed, she might have the complete pipeline of knowledge and sensitivities needed for this work. Some additional names in this area -at the intersections of social science, administrative responsibility, data mining for attributed graph analysis, etc - that could be useful but might not be as good of a fit:
  - Leman Akoglu : her group (Heinz, [The DATA Lab](#)) certainly has sufficient technical expertise, seeming to have focuses that match the algorithmic aspects of the problem being proposed here. I believe she has done some social-analysis work - however, it is not clear to me if she or her group has experience working in DEI or algorithmic bias /fairness issues and the numerous issues that emerge in those sub-areas.
  - There are a number of people at CMU who focus, at least in part, on networks, network analysis, and data mining more generally. One that comes immediately to mind is Christos Faloutsos, who naturally is a wealth of expertise. I am not sure of his focus on social problems, however. Moreover, while all the people I list here are busy, Dr. Faloutsos strikes me as already having a fully committed schedule - as such, he might be able to provide some suggestions and do some very high-level advising, but I doubt he personally could get involved in the effort. *That said*, maybe one of his students could help with the work - he does have some people working on countering human trafficking ; this, while a largely different set of techniques, does indicate some concern for societal issues in his group. Further, Dr. Faloutsos could probably provide some good suggestions of who else to contact in CMU.
- Some names in Deitrich's statistics department of people who *might* be interested in joining this work (or at least could have suggestions). I am not as confident on these fits, but it seems nice to mention them to help round-out this list - everyone else mentioned has been from SCS and Heinz, I would feel silly not to point out that Deitrich also has relevant resources.
  - Brian Junker (faculty, andrew ID: bj20)
  - Nynke Niezink (faculty, andrew ID: nniezink)
  - Xiaoyi Yang (PhD student, andrewID: xiaoyiy)
  - Kayla Frisoli (PhD student, andrewID: kfrisoli)

#### A Note on Relevant Resources at CMU Outside SCS:

Note: I say the below just to give an idea of what resources are available and be clear that they are not limited to SCS, at least in respect to qualifications.

Truthfully, there are many people - particularly faculty - at CMU who have a skill set and background that may be reasonable matches, and certainly so if consulted as part of a team (not to say we necessarily want more than one person, but just to indicate that there are many ways to get the needed expertise together). While SCS seems like the first-stop to find someone with the necessary technical background, in reality it looks like both Heinz and Deitrich have focuses with similarly good candidates, *some of whom do not (to my knowledge) have joint appointments in SCS*. As a general idea of the expertise in these other CMU schools: Heinz has a PIT (Public Interest in Technology) initiative and a whole focus on information systems and management (which does range from social to technical considerations), while Deitrich has whole departments for Economics, Information Systems, Social and Decision Sciences, and Statistics. While not every practitioner in these arenas has the computer science expertise desired to oversee the analysis, many do - for example, there are many social scientists and economists that do the sort of

network analysis in question, but certainly not all. The point is this: SCS is not the only resource available to us here. In fact, in order to further avoid biases in the analysis, pulling from outside the school *might* be advisable - maybe.

I would be happy to discuss, expand any of my points here, or answer any questions these points evoke.

-yours  
D Bayani

### 3 Briefly: A Few Points of Considerations for Discussing ML's Impact and Actions CMU Can Take

This committee had extensive discussion at least twice (that I am aware of) in regard to promoting responsible and ethical use of ML in CMU SCS. The PhD Anti-Racism committee has also discussed this subject. I believe that these bodies definitely contain individuals that are knowledgeable of several questionable ML applications, and have proposed several actions to limit CMU's role. I had the privilege to be part of several discussions on the subject in these two committees. Here, I would like to briefly expound several points that may be important and that, as far as I can tell, are not often raised by others in these discussions. Having been explicitly asked to speak about this subject before, particularly [REDACTED] it seems reasonable to provide these comments for reference during future MAC discussions.

To begin with, actions that are typically proposed should be mentioned. Two undertakings that are often brought up are (1) requiring members of SCS to take some ethics courses prior to graduating, (2) have SCS limit its involvement in (USA based) military and predictive policing projects. An additional third, to complement (1), is the idea of incorporating related materials into already existing classes. These three proposals have pros and cons, and may be contentious (particularly (2) ), but they at least seem like reasonable actions<sup>4</sup> that could have non-negligible effects.

With that noted, below I attempt to promote greater explication and awareness of several high-level planning considerations:

#### 3.1 The Importance of Incorporating Awareness of Uses and Responsibility into the SCS Ethos

I believe an essential high-level goal is to make these impact and ethics considerations part of SCS's ethos; this is in contrast to seeing it as secondary to immediate action on the local scale, or simply a side-benefit. The hope is that a proper spirit will breed action beyond the borders of the here and the now. I am not objecting to acting now; my point is that our ability to act appropriately in the future is dependent on our ability to develop people now that will later act appropriately.<sup>5</sup>

People entering CMU should get the sense that members of SCS are well aware of what sort of serious problems can occur; people exiting the school should have some concrete understanding of societally high impact and presumably unintended ways the technology they develop can be used.<sup>6</sup> Naturally, different people will absorb such an environment differently, but the fact that some may walk through unchanged should not discourage adoption ; not everyone in a class gets an A, and (in most schools) more than a few people fail a course - but the fact that less than a majority earn A's is hardly ever taken as a reason stop teaching a course outright.

The reasons to promote this ethos are based on the lack of the current awareness inside SCS, and the fact

---

<sup>4</sup>Here, "Reasonable actions" in that "reasonable people" would likely find there to be sufficient ground for a non-trivial discussion / debate on the subject. This is not meant to imply my support or lack-there-of.

<sup>5</sup> By analogy: fire departments not only put out fires when they occur, but also invest in fire-safety infrastructure and education; further, most cities do not disband the fire department as soon as the current blaze is extinguished — consider how awful future incidents would be if no firefighters existed at the time of occurrence. While this is a strong example of consciously supporting both preparedness and timely action, it problematically centers around fire-fighters, whom are professionals or highly-trained volunteers. It is not feasible — or even necessarily desirable— to have all members of SCS become "complete experts" in regard to ML's societal impacts. A better analogy to follow might be that of learning CPR, which many people in a variety of roles are taught (ambulance drivers, life guards at pools, ship captains, etc.): points of interest are that: (a) one does not need to be a doctor to use CPR, (b) it is taught to people who will use it in the future, if ever at all, (c) learning CPR is considered part of the required training for an occupation that might otherwise have little to do with the medical field.

<sup>6</sup>I am not sure SCS members are so naïve, but there are pundits whose perceptions are that computer scientists are short-sighted or, worse, small-minded. For example, there was a discussion on NPR's All Things Considered where reports claimed — in the context of Facebook's apparent "connecting people is inherently good" philosophy — that many people producing technologies have childish world views (I am having difficulty finding the correct reference to cite, but the program was particularly explicit in making this claim). While the judgment of others is something that should worry us little, we should hope/aim that our thoughts and actions clearly demonstrate their judgments to be false.

that these concerns are neither limited to the present nor to CMU. Being more verbose, three clear reasons are:

1. At current, it is primarily people who really care about “responsible” ML-use that hear much about it — for the rest, the point is abstract and vague. It is questionable whether most would be sufficiently aware of probable consequences prior to deployment or development of techniques. It is certainly unfair to expect people to see the future, but we should do due-diligence in ensuring others have the full perspective necessary to do their own due-diligence.
2. As already stated, we want SCS and students produced by it to be prepared for the future. Issues relating to abuse of ML are not going away and will take radically different shapes and scales over time. As such, there needs to be a community to actively monitor, be wary, and adapt as uses evolve — not simply throw bandages on general problems when they mature enough to be obvious and large.
3. The geographic and institutional scope of CMU’s impact extends well outside the university. The works that CMU directly gets its hands dirty with, while important, are only a small part of the picture. Alumni go-on to many places and have long careers. Publications spread research globally. As such, the impact of CMU’s technology — and the abuse of ML more generally — has state, national, and *international* impact. It is important to emphasize that this is not hypothetical, but backed by observation. This university is not the only actor contributing to the growth of ML and AI more generally, but that does not relieve us of the duty to be conscious of where our works might be deployed and how.

Of these three points, I believe the first two are already understood by the committee, though implicitly so. The third, as far as I am aware, it not a item that is often considered. For instance, there was discussion concerning CMU’s involvement in USA-based predictive policing and military efforts: most suggestions for actions simply called for limiting (current) CMU involvement or funding. These suggestions, while actionable, did not appear to consider the larger scope. It is possible that the scope of those discussions was limited out of practical concerns (i.e., it is easier to tackle the concrete, local and immediate then the abstract, global and eternally<sup>7</sup>) — indeed, if the school started “promoting the right ethos” but did not review these contentious projects, it would reasonably be seen as a cop-out. That said, however, I do not recall hearing larger boundaries being discussed (particularly across nations) except for when I raised the issue. It is possible that many automatically assumed a broad scope, but at the very least it sounded like some did not consider the impacts beyond (a) the projects CMU is currently a part of, and/or (b) the United States. If I am lucky, I sound foolish here by stating points that are already well understood. Even if so, I do believe there is often value in making lines of consideration and scope explicit — if, that is, they have not already been made explicit somewhere.

A closing thought here on the role of the school ethos:

If we want to take an optimistic stance, it is sometimes true that thoughts and attitudes are contagious. One hopes that if CMU sets an example, it will be picked-up by members of the academic and industrial communities outside our alumni. In short, an ethos can spread in time, space, and across people. However, the limitation of promoting “a proper spirit” is that the notion is soft and “results” hard to measure — further, it complements, but does not substitute for, acting in the present on the materials the university has immediate access to (i.e., stopping involvement in problematic project X). Indeed, one can have the “proper spirit”, but if it results in exactly the same actions developing, then there is no utility in this context.

### 3.1.1 Use of Examples or Case-Studies

MAC has discussed backing the addition of technology-and-social-good-related classes to the required curriculum. This is reasonable. There has also been discussion of adding content to existing classes. Toward the latter suggestion of adding content, I digress briefly here on the power of examples for learning these soft concepts.

---

<sup>7</sup>Case in point: “promoting world peace by promoting the spirit of peace globally” is rarely considered useful, particularly if it is not accompanied by other actions.

Educating with case-studies may help more people in SCS have their senses heightened towards these issues, and may help lay a common ground-work for discussion. There is already extensive precedent for case-studies in CS, for example in studying software security and information systems<sup>8</sup>. Further, examples are very powerful ways of helping vague concepts be clarified, absorbed, and remembered. The ways in which lawyers and philosophers go about their business comes to mind. For people outside those disciplines (such as myself): Think of what lessons from history you know — think about opinions you have about government, justice, and other issues — what proportion of your *strong* feelings were shaped or emboldened by considering specific examples? You may have learned things from historical examples without ever having had put words to the lesson.

When dealing with abstract notions, examples often help individuals gain a more precise picture on where they, personally, view the line between acceptable and unacceptable to be. Further still, while overcoming the sway of personal biases in nuanced debate may be an insurmountable challenge, having solid situations to illustrate a stance's limits may assist in evening-out perspectives.<sup>9</sup>

Under a certain analysis, there might not be much of a point in this sub-sub-section — taken too liberally, the point here is that “examples” are useful education tools — but if one doesn't use examples, then all one has is theory, and I am not rejecting the usefulness of wrangling abstract concepts — so it may sound almost like, in the toolchest of human problem-solving strategies, I am praising one half of the tools and claiming that their utility is non-obvious. My writing here, then, would show me to be a self-important imbecile, something which hopefully is not being demonstrated.<sup>10</sup> The point has been to discuss examples / case-studies as added elements to existing courses in order to illustrate ML-and-Society issues; focus has been on the specific utility of the approach. While kept implicit thus far, a point of this discussion has been the bang-for-the-buck case-studies might have. In particular, I doubt much room can be made in existing classes for added ML-and-Society content without radical restructuring. To achieve the desired goals, it might be most prudent to spend that time on concrete illustrations as opposed to providing theory and hoping that students connect-the-dots regarding impacts (a slide with citations for theory papers is always possible for interested students).

In regard to specific examples to discuss:

To remain fair-handed while avoiding too far of a digression, I spare listing specific examples here (I would be happy to discuss cases in a meeting). However, I encourage the reader to visit some collections of problematic examples (e.g., these: 1, 2, 3, 4, 5, 6, ) then follow-up by searching for related stories on at least two reputable news sites, such as BBC, Al Jazeera, Voice of America, etc. I would suggest against relying on one news source or news sources located in one nation; in particular, it is a bad idea to rely only on state-run media.

Currently in the USA and EU, there is focus and concern about the privacy<sup>11</sup> and public-discourse impacts of industrial deployments of ML. Additionally, in a broader context, I have concern over ML-assisted surveillance enabling censorship and targeting — physically, in the real world — of dissenting individuals/groups.

In regard to a few implementation considerations:

---

<sup>8</sup> Arguably also in ML when discussing specific papers and algorithms -there is a spectrum.

<sup>9</sup> This point specifically came up when discussing DEI-related suggestions for SCS administration. There was a proposal to limit police presence at CMU gatherings. More than a few people found the proposal problematic as it was phrased, appearing to be too far sweeping by mostly forbidding police attendance. As one commenter put it, the discussion needed to be more nuanced and the proposal as-written “threw the baby out with the bath-water”; there seemed to be consensus in the MAC about this general point during meetings prior. As I commented then:

To help shape this discussion, it would also be helpful to have particular examples of concern that help inspire this ask. Some good examples of use and some bad examples of use would help build understanding. Unfortunately, police use in most capacities is one that can quickly go to straight to a person's general leanings - e.g., “police make me feel safe” v.s. “police make me feel stifled and unsafe”. In abstract, I suspect those feelings dominate in a otherwise (potentially) nuanced discussion.

I suspect such a pattern will not be uncommon in controversial subjects.

<sup>10</sup> I am a fool (and a tool), but with luck it is for slightly different reasons.... Does a fool know why they're foolish?

<sup>11</sup> As far as I am aware, in principle there are more restrictions on what the US government can do with your data than what these companies can do.



██████ suggested during the MAC meeting on September 20th, 2020 to have a curated list of examples added to syllabus. This appears to be a reasonable idea. Not only does it reduce instructor work-burdens, but it also helps ensure satisfactory discussion, since we do not expect all educators in SCS/MLD to be equally well versed or sensitive to these issues. It remains to be discussed whether there should be mandated, stock examples to cover or simply a list of examples that are encouraged. On one hand, a mandate helps edge against insensitive instructors, but at the same time, we would expect a subject-area professor to be a chief expert in regards to which examples are most relevant and impactful. If the process to approve a professor as a course instructor is non-trivial, then the best of both worlds might be achieved by having the department sign-off on the examples the professor plans to use. If the review process is lightweight or very granular, than providing strong recommendations may be the only option, short of a more radical change in department practices and culture.

### 3.2 Educating While Attempting Not to Act as Ethics Police (If Deemed Desirable to Not Do So)

To be clear: we can try to separate discussions of what is “right” or “wrong” from providing examples of societal impacts. An event and value judgments about an event may reasonably be said to be separate things.<sup>12</sup> We can educate about examples of use with clear, direct social impact without necessarily giving clear social judgments; some tone of approval or disapproval might seep-in inevitably, but certainly the degree can be limited — kept away from being overt.

This is an important aspect to keep in mind since different people — especially coming from the diverse experiences that a globe has to offer — have different value judgments for specific events. Having SCS stray away from value judgments helps prevent CMU from acting like a broad ethical authority, which is desirable since encouraging the opposite: (1) is, in-and-of itself, questionably appropriate<sup>13</sup> (2) may turn off students<sup>14</sup> (3) may serve to silence minority / rare perspectives that simply have not yet been considered or appreciated by the wider SCS. I am *not* necessarily saying CMU should *stay away* from issuing value judgments, but to be clear that promoting an atmosphere of impact education and awareness does not necessarily require CMU to commit to being an ethics judge on all subjects; maybe SCS wants to be an overt ethics judge, but that would not logically follow from simply wanting to educate.<sup>15</sup> It might be naïve to suppose those two things can be separated completely in practice, but at very least they can be distanced in principle. In addition to these concerns, there are more immediately practical implications of the opposite; it can be very hard to gather administrative or community support for initiatives that appear to involve one person shoving their opinion down another person’s throat.

Inside of CMU, there might be precedent for walking this line: for example, University Health Services (UHS) offers a variety of contraception options and does make efforts to educate undergrads (particularly

---

<sup>12</sup>At least in common understanding. I believe some argue that they are both inseparable parts of the same picture, or at least that one without the other provides little value. The discussion spans ontological and pragmatic concerns. Philosophical debate and nuances, however, should not totally block consideration of the statement.

<sup>13</sup> The mess resulting from the statement “It is unethical to act as an ethical authority” is not totally lost on me. A less questionable stance may be “it is inappropriate to project personal values onto students as though they are objective fact”; it remains to be precise in how that is not simply different verbiage for the same troublesome self-reference as the previous sentence. Really, this “trouble” is most likely a result of muddling judgments about different entities or failing to curtail scope of authority. If we insist on resembling earlier phrasing while attempting to raise fewer alarms, one would be better to adopt “It is unethical to act as an ethical authority on issues outside of scope X”. To avoid an all-too-easy long digression, I leave this subject thread here, despite the lingering issues of precision in meaning ( and moreover how silly I look from this immature subject treatment ). If nothing else, it should be clear from this amateurish and ridiculous writing that one would not want me to be the judge.

<sup>14</sup>or earn the ire of corporate sponsors and administrators of various regimes

<sup>15</sup> I remind the reader that the context in discussion is course-content, etc. Earlier in this document was discussion of DEI issues - in that context it should be clear that we do want CMU SCS to discourage certain kinds of behavior, setting both practical and ethical standards.

Digressing briefly: hopefully context, scope, and specific circumstances help illuminate the proper, immediate actions and division gradients here (division gradient as opposed to division line). In the broadest conception, deep issues may emerge. I believe Bertrand Russell pointed out issues with the phrasing “I am intolerant towards intolerance”, and the notion of “tolerating intolerance” is clearly unappealing (by the way, who is judging what is or is not tolerant?). In any case, it should be clear that I am trying not to paint with too broad of a brush.

new freshman) about the existence of those services. In that case, it appears reasonable to say CMU is willing to offer services and educate about their existence without necessarily encouraging or discouraging the practices that lead to them. Admittedly, the philosophy behind UHS's actions may be based on the harder-to-borrow grounds of "harm minimization"—under this school of thought, if people are going to do heroin, and you have minimal control on curbing their usage, it is better if you at least provide fresh syringes so to minimize their chances of contracting blood-based diseases from dirty needles. As such, the philosophy guiding those actions may have limited applicability to SCS's educational practices, due to (a) the differences in risk/benefit types and (b) the degree/direction that SCS would (potentially) like to influence behaviors / perspectives. At a minimum, though, a few lessons are likely transferable to SCS regardless of whether the same exact justifications for initial actions are transferable.

This entire discussion relates to debates on what it means — or whether it is even possible— for news papers, encyclopedias, history books, etc., to be objective. Editors, writers, and reports select a subset of available information to present and make decisions about the nature of presentation,<sup>16</sup> and such content decisions impact how people will perceive the item discussed. Regardless of whether or not it is possible to provide content without subjective influences, we can at least act (1) in good faith and (2) to avoid many forms of overtly and egregiously one-sided discussions. If we go forward with educating, mistakes almost certainly will be made, then corrected - then new mistakes will appear; one would hope that the cycle approaches convergence to a reasonable, time-varying balance. One would hope that the SCS community can tolerate imperfect discussions so long as we honestly pursue "good"/"healthy" discussions.

---

<sup>16</sup>For instance, the order facts are delivered - it is often hard to undue the influence of people's first impressions, so order and emphasis of presentation are not necessarily trivial from a "psychological" point of view.

## 4 In Regard to Archiving MAC Materials

### 4.1 Suggestion for Next Iteration of Improving Archiving

To get to the chase, I list in this subsection thoughts on how we could immediately improve archiving; I spare other options and some further rationale for later.

While not perfect, using GitHub to collect and securely store materials - while allowing easy transportation of the associated records and metadata - does not seem like a bad choice. The committee members who oversee the Google Drive and/or the MAC website could act as the administrator. The archive could be hosted in the same GitHub account as the MAC website. The administrator could oversee the master branch<sup>17</sup>, and each member could request material be added to the record through merge-requests. The repo should not be public — it hampers discussions and internal communication otherwise; the MAC retaining non-public records is not without precedent, since the Google Drive, Slack page, and other committee communication infrastructure requires credentials beyond simply CMU Andrew credentials.

This approach does have downsides. For one thing, it does not make archiving automatic, but something that individuals have to do. This has the potential pro that only desired statements are archived, which is potentially useful if discussion of a controversial topic where people may feel sensitive having preliminary comments recorded. The associated con, however, is that it is less likely material will be archived due to the extra effort required to remember and perform the backups. Hopefully, those who care enough to produce material care enough to push it to GitHub after — but even if the likelihood of voluntarily archiving is higher for such people, it is still not certain, and thus more likely to fail in some cases. This said, most MAC members probably already know how to use git (unlike many other non-automatic platforms we might otherwise adopt), and thus the barrier of a learning curve is minimal.

Imperfections aside, it seems imprudent to let the perfect be the enemy of the good here. There appears to likely be reason, given the minimal infrastructure required, to make this option available for the committee. If better ideas come later, the MAC can adopt them - there is nothing to stop us from moving to another platform, particularly given how Git facilitates moving data and metadata. While difficulties later-on are possible — for example, if we later adopt a platform that makes it difficult to import git data — they seem no more likely as a result of using git now than any other platform. Even in the worse-case where migrating proves burdensome, at a minimum a git-archive can be compressed and stored on a new system as a large binary. As with any infrastructure, we would just need to ensure committee members are aware of the resource, and remind them to please use it in cases where they are comfortable recording the information.

A note on previous discussions of this subject: to the best of my knowledge, chats using the explicit term “archiving” started near the beginning of September - the first discussion I am aware of was between [REDACTED] and myself when [REDACTED].

[REDACTED]. However, when I raised this matter with the committee on the following week, it appears some related thoughts might have circulated previously. In particular [REDACTED] mentioned that some aspects of data retention and exporting were considered when trying to select what messaging platform to use for the committee. I am not quite sure what the discussions and decisions were in those previous conversations. As such, I might be rehashing points previously considered by the committee but that were not widely communicated; I beg pardon if so.

### 4.2 In Regard to Privacy and Consent to Sharing Other People’s Contributions

Archiving committee material becomes more complex in the face of trying to respect committee members’ privacy and agency over their contributions. As an example of undesirable occurrences: suppose the committee discusses in writing a highly controversial topic and someone states an opinion they would not want shared more broadly - or a statement whose wording fails to express the author’s original intent<sup>18</sup>.

<sup>17</sup>I particularly suggest enabling branch protection, like preventing rebasing in merges, etc.

<sup>18</sup>This includes, for instance, an statement being taken out of context due to only a strict subset of written exchanges being archived.

This is a concern to balance against the clear benefits of archiving the committee's collection of information resources, discussion, debates, and history of actions.

It seems likely that the best way to address these privacy and agency concerns are to make some sort of addendum to the MAC constitution or create other solid committee "legislation"<sup>19</sup>. The MAC might not be formally required to make such a statement, but being straight-forward and up-front about how contributions will be retained and shared is the right thing to do, external requirements aside. I suspect a good model would follow the spirit of: (1) text communications are opt-out (with revisions allowed after-the fact, as detailed below), (2) voice, images and video are opt-in.

For example, in regard to text being opt-out, the MAC might decide that all X content from source Y is archived, but that members can request anonymization or redactions (but not arbitrary edits) after data collection. Bear in mind that data goes through two different phases after creation: first it is collected and then, later, it is archived. It would seem reasonable to allow the following:

**Requests that Members May Make:**

- Option A: material they themselves produced be left out entirely (if not yet archived)
  - *A few points to consider:* It is debatable what should be done if material from a person has already been archived. Likely, the circumstances play a role: if a person explicitly agreed to have something archived then changed their mind, it seems different than if the process occurred without their knowledge or consent. That said, if archiving occurs routinely, and if the point of this discussion is that archiving content is opt-out, not opt-in, then it is not clear that a committee member could "fail to know or consent" in a meaningful sense prior to archiving occurring (unless they are new to the committee and thus unaware of the policy). That said, if archiving occurs more quickly than a member can make an edit request, it would seem unfair to deny their request. Maybe we can do the following:
    - \* after content creation, at least two weeks will go by before archiving occurs
    - \* if a member wants a content-edit after archiving, a simple majority of the MAC must approve via email (not by meeting because it is hard to ensure a simple majority at the meeting where the issue might be raised). Obviously, an edit-approval email would either be anonymized or not recorded — otherwise it may defeat the point of the edit in the first place.
- Option B: material they themselves produced or directly contributed-to be anonymized with their choice of the following:
  - Option B.1: *Identifying information be stripped.* For example, the name becomes "person1". If there are three people who want to be anonymous, they will be recorded as "person1", "person2", "person3", for instance. This numbering (as opposed to referring to all anonymous people as "personX") helps ensure references to different people are kept straight, and be clear that different entities might believe different things. See section 2.1 for an example. While this stripping-approach removes the immediate identifying information, it still allows inference of who-said-what by process of elimination and writing patterns. However, the likelihood seems low of future MAC members being able to infer authorship without expending substantial effort, especially if they act appropriately<sup>20</sup>; that is, the concern over de-anonymizing decreases as time progresses, both because of behavioral expectations and decreasing familiarity with the previous membership rosters.
  - Option B.2: *Provided contributions are removed and clearly marked summary text is inserted instead, in addition to Option B.1* . This reduces risk of easy deanonymization. However, it increases the burden on the committee. As such, the people requesting this should understand

<sup>19</sup> What is other verbiage here? Committee "policy"?

<sup>20</sup> "Appropriately" in the sense that a reasonable person would agree the member acted in accordance with the spirit of the committee's policies. Operationally, this is determined by the committee debating then voting on whether a member acted in violation of the principles in question.

that the replacement text might be short, and might not capture the full, original meaning. Recall that removing content is allowed (see Option A); the distinction here (Option B.2) is to provide some idea of what the original content was for future readers. See the text in section 2.1 that looks like “[summary: [...]]” as an example where I attempted this.

- Option C: have a clarifying footnote appended to material implying their attitude on a subject:
  - Elaborating on this option : we would add a note to the material in question indicating the chunk of content for which the member (either specified as an individual or left as an unnamed committee member) did not necessarily subscribe to. For instance, an asterisk or a footnote with line numbers might state “At least one committee member requested that it be noted that they do not necessarily agree nor disagree with this statement”. That’s a verbose and polite way of saying “don’t think everyone who attended this meeting agreed with this statement, even if the committee passed it — and someone felt strongly enough to have their disagreement noted”. Again, this is useful for controversial topics or other cases where a member feels inappropriately lumped-in to something they dislike.

Above, notice that the outlined focuses above on (1) giving people control over their contributions and (2) tries to prevent “information leak” of *their* opinions or thoughts that they might feel particularly sensitive about. However, MAC members won’t have the right to, say, edit the content of an email if they only happened to be a receiver of the email. Obviously, there are grey-areas — for instance, if an email comments on the person specifically — but the spirit of the statement here should be clear: generally speaking, the more strongly content implies something about you or your thoughts, the more leverage you should have convincing the committee to anonymize or redact it.

Summarizing the points thus far: for written communications, adopting an opt-out approach appears reasonable so long as we are explicit about it up-front and allow for very robust cleaning of personally concerning materials.

For real-time communications, such as Zoom meetings, live spoken word, or other non-pre-recorded media, it seems best to apply an opt-in policy: it is often hard to find the right words and easy to misspeak, especially when discussing difficult subjects<sup>21</sup>. In other words, opt-in seems the most fair for cases where members do not have much time to consider their actions or reflect on phrasing.

A closing note for this part of the thread on privacy issues and addressing them: Again, this all looks OK prior to material being archived, but issues emerge if these edits are allowed after archiving. For example, this opens the door to historical revisionism<sup>22</sup>; how we would avoid that and whether we want to prevent it is a remaining question. One hopes that entertaining these worries now is getting far ahead of ourselves. With luck, the MAC might only tackle such a concern once a decade.... in which case, the specifics of the situation will hopefully make the decision clear. In other words, it hopefully suffices to say these are bridges to cross once we arrive at them.

A bit of a digression to provide clarification and highlight several tradeoffs can be found in section 4.4. The content of section 4.4 was not included here due to length, subject, and writing quality considerations.

## 4.3 Other Implementations Details and Thoughts

### 4.3.1 Some High-Level Factors to Consider

The method of archiving is a balancing act of at least the following items:

---

<sup>21</sup> Not to mention that opening this avenue for criticism allows originally innocent statements — such as a joke that meant little in context — to be subject to much scrutiny. No one wants to hamper engagement by killing the mood and spirit. Naturally, I say this under the assumption that all content was appropriate in context — I **do not** mean to imply jokes of bad taste are OK so long as the immediate audience is unoffended. The committee has thus-far done a good job keeping humor appropriate and of minimal bandwidth consumption, as one would expect from mature people. (not to imply this document’s author is mature by any standard).

<sup>22</sup>For example, someone may claim they did not support something that they later grew ashamed of when, at the time, they really did.

- durability : we want to make sure data, edit histories, metadata, etc., are retained over the long term and cannot easily be made unreadable (i.e., deleted, corrupted, etc.)
- easy of use : if it takes serious effort to archive anything, then, realistically, it will rarely ever happen
- clarity where information should go and will go : this relates to ease of use , privacy, and avoiding confusion about how to best communicate
- practicality in implementing the scheme : we don't want to spend too much time fantasizing unrealistically over possible implementations
- ensuring members are comfortable : it is important to ensure that recording does not (1) prevent honest discussion of sensitive subjects, (2) make people feel self-conscious speaking, or (3) create other unintended barriers to effective and genuine interaction.

### 4.3.2 The CMU MAC Mail Server

Having checked roughly two weeks ago (while trying to find copies of the emails that are copied in this document, in fact), it appears that the MAC mail server has a hook for archiving messages sent to it (see SCS MAC Mailserver Login). However, when I looked, it was empty. Touching base with ██████████ ██████████ stated that, to ██████████ knowledge, no SCS mail server actually archives the content. However, ██████████ suggested I touch base with the SCS Help Desk if I really wanted to check.

### 4.3.3 Backing Up Slack Content

Miscellaneous thoughts on backing up content from Slack:

Slack does allow users to export their data. However, for non-commercial users, it does not provided the files that were uploaded, only file names as placeholders. See:

Slack:Export Your Workspace Data

and

Guide to Slack Import and Export Tools

The MAC Mail Server has a setting where it can receive mail that is not forwarded to members of the mailing list. If the server archiving was functional, it might have been possible to simply have Slack messages sent to the mail server and archived there.

Having an —unrealistic — pipe-dream for a moment: we could write a Slack bot to automatically push content to a MAC GitHub via some webhook. Sure, its overkill, but we're CMU SCS students - do we really need an excuse to build (ro)bots? The process could be "as simple" as what follows:

```
snapshotTime="300"; # every 5 minutes, add changes. Can be adjusted
# as we see fit
botName="nameOfThisCommunicationChannel_gitHookBot";
while(True):
    git add ./content/* ;
    git commit -m $botName" : Autocommit"; # automatically
    # does nothing if there are no changes to add.
    GIT_SSH_COMMAND="ssh -i pathToGitSSHKey" git push; # better
    # handle any errors at this stage. Also, only run it if origin
    # is out of date.
    sleep $snapshotTime ;
```

See

Slack: Create a Bot for Your Workspace

for further details. When I looked, it did not seem like Slack would easily support this, but it might be nice if I am wrong.



#### 4.3.4 Google Drives

Google Drives is great for synchronous collaboration, providing a shared editing environment, real-time data entry/editing, and an easily accessible space for files. General downsides of the platform, however, are that it is hard to export revision history, hard to export metadata, and not durable - or at least not as durable as Git very easily can be. Further, at the moment, the MAC has set up the Google Drive to allow any MAC member full access and modification privileges - in the context of easy collaboration with trustworthy committee members, this makes sense on the platform, but it has potentially serious trade-offs.

#### 4.3.5 Size Limits in GitHub

GitHub lets you push arbitrary binaries (pdfs, txt files, html, images, etc.), it just complains whenever a single push is too large. I seem to recall “vanilla” GitHub also enforcing file-size limits (something like 100MB for the largest file), but I have personally gotten around that in the past using the Linux utility “split” — as such, the only concern is the commit size. Another capacity limit is the size of the repo - looking just a week or two ago, it appeared that GitHub has a limit of 1GB for repos. This is all in reference to GitHub without using the LFS (Large File Storage) capabilities: LFS alleviates all or most of these concerns. More on GitHub LFS can be found at:

Configuring Git Large File Storage

GitHub LFS: About Storage and Bandwidth Usage

About GitHub LFS

Most likely none of these limits will prove to be an issue for the committee, they’re just good to be aware of. We’ll handle sharing files on the order of hundreds of megabytes if it ever comes up - most likely, it won’t be an issue except once or twice a year. Most of the committee’s communications have been well facilitated by texts and an occasional PDF. If we ever start sharing videos or images related to our endeavors, capacity limits may become an issue. In such a case, however, I suspect GitHub LFS is sufficient with its 1GB limit on storage and monthly bandwidth, especially if content is compressed prior to committing.

#### 4.4 Addendum on October 4th, 2020: Clarification of My Actions and Stance, and Highlighting Several Tradeoffs

In at least two paragraphs of this document, I have mentioned names of committee members and a small description of their contributions. The motivation was to share credit as appropriate and attempt to recognize valuable contributions. Unfortunately, I might not have done that enough in this document, as I leaned towards avoiding misrepresentation of other’s affairs and protecting individuals’ potential preferences, as outlined previously; indeed, I had initially referred to [REDACTED] as “a committee member” to avoid this issue. To be clear, however, when other parties (though perhaps anonymous) contributed to something, it has been noted; attempts have been made to recognize contributions while not running afoul of individuals’ preferences.

The question from this, then, is why I felt comfortable mentioning people in the cases I have. I suspect that the average reader might not wonder this since, with luck, the previous mentionings were innocuous. It is worth being explicit about the consideration, however, since it illustrates a trade-off between risking respect for an individual’s privacy preferences and providing them proper credit for their contributions. I decided to list names where I did for at least the following reasons:

1. There was an opportunity to provide acknowledgment where it was earned.
  2. The content referred to was stated in MAC meetings. As such, the content was reasonably expected to be shared with the same audience as this document.
- It is not necessarily the case that limiting my notes to the same audience made them permissible; in general, a commenter may make statements to an audience expecting the words to be temporary, not recorded in any way. In the MAC meetings, however, we do take meeting minutes, which are a written record paraphrasing the conversations with large granularity.

- That said, if I am suggesting that meeting minutes have yet to be “archived”, then it would seem implicit consent is not yet granted for the further step of archiving. That is, when speaking in a meeting, there might not have been an expectation of durable, long-term records.

My comments here are going to extremes, but the point is that this justification, while potent, might not be bullet-proof.

3. I paraphrased their contributions at a high level of granularity. I did not share quotes or material they themselves produced.<sup>23</sup>
4. The people named were sharing material of a factual or informative nature, etc.; I believe that a majority of reasonable people would agree<sup>24</sup> that the material was sufficiently uninformative in regards to the opinions, beliefs, stances, etc., of the named-persons. The division is not solid, of course. For example, you can guess the political affiliation of someone based on what facts they choose to share or what actions they suggest. That noted, I do believe most reasonable people would agree I am on the acceptable side of the line. Further, the highly-granular paraphrasing I performed helps hide elements of the original speaker’s belief.

Obviously, if any of the people I mentioned would prefer, I would happily anonymize their names upon request, using the options listed in section 4.2.

I might have made the wrong decision in some of these instances, particularly since I did not request permission first. In the worst-case, maybe I was lazy, maybe I feared ridicule for making a big deal of something trivial, and maybe my considerations are just rationalizations; I hope not, but maybe that was the psychology. In any case, I might come back and redact later, if further considerations show it as better. In general, I am not in favor of unredacting (in this context, it would defeat the point), so if I remove the material, I won’t put it back in later.

Truthfully, I suspect the committee — to make life easier while governing a *student* organization of small size — will simply adopt some granular policy that addresses the area in general terms as opposed to trying to fantasize about cases and sub-cases galore. The principles are valuable, but I do recognize that we are not a government and that actual work may be impeded by pretending too much to be ministers. I support a more granular approach, as I outlined, but I recognize the fact that this may be seen as childish, getting ahead of ourselves, a waste of time, or simply unnecessary. With luck, however, the high-level considerations that motivated this rambling will help clarify issues to consider when the committee decides to take a stance.

A blanket statement that is a bit vague but hopefully still serves to rule-out many undesirable stances: It is appropriate to describe the contributions from other people even if done in a private channel; in the case of private, person-to-person communications, the debatable aspect is whether (a) their names should be shared without permission or (b) it should be simply noted that another party (e.g., person1, person2, etc.) contributed in particular fashions. Whatever the channel, it should be clear that when other people provide value to an endeavor, reports should not be stated as though the report author was the sole contributor. A report with anonymized contributions should provide enough room and clarity so that most reasonable people would believe it is not inconsistent with the real, de-anonymized story. I believe everyone on the MAC would agree with this paragraph; the reason for this statement was to highlight at least one high-level, agreeable and easily implemented stance which the committee might decide to then refine, expand, and adopt. Prior points, while worthy of consideration, have neither been as concise nor as obvious, and thus have higher inertia as written.

---

<sup>23</sup>Division lines can be elaborated further - the ideas in copyright law about derived materials might serve as inspiration. I am not sure plagiarism considerations directly apply here: the question is not whether I am taking their content without due credit, the question is whether I have used enough derived content (while citing them) to cross a threshold of concern.

<sup>24</sup> Notice that the standard I am aiming for is not an individual’s judgment, but the judgment of “most reasonable people”. Unfortunately, such a jury is usually not available, so one has to mentally model them - hence the part saying “I believe...”. Again, the point is not what my personal stance is, but the closest I can estimate how most sensible people would rule.

## **5 File Size and Hash**

This file's sha512 hash and size were communicated over email to [masters-advisory-committee@cs.cmu.edu](mailto:masters-advisory-committee@cs.cmu.edu) on day5 month10 year2020 (UTC-0 time); comparison to the values reported there should help ensure this document has not been modified in an unexpected or unreported manner.